#109

An Interview with

ROBERT E. KAHN

OH 192

Conducted by Judy O'Neill

on

24 April 1990

Reston, VA

Charles Babbage Institute
Center for the History of Information Processing
University of Minnesota, Minneapolis

Robert E. Kahn Interview
24 April 1990

Abstract

The interview focuses on Kahn's role in the development of computer networking from 1967 through the early 1980s. Beginning with his work at Bolt Beranek and Newman (BBN), Kahn discusses his involvement as the ARPANET proposal was being written, his decision to become active in the implementation of the network, and his role in the public demonstration of the ARPANET. The interview continues into Kahn's involvement with networking when he moves to IPTO in 1972, where he was responsible for the administrative and technical evolution of the ARPANET, including programs in packet radio, the development of a new network protocol (TCP/IP), and the switch to TCP/IP to connect multiple networks. This interview was recorded as part of a research project on the influence of the Defense Advanced Research Projects Agency (DARPA) on the development of computer science in the United States.

ROBERT E. KAHN INTERVIEW

DATE: 24 April 1990 INTERVIEWER: Judy O'Neill

LOCATION: Reston, VA

O'NEILL: I want, as I mentioned, to focus on your network experiences. You covered some of that in your previous interview. You described in your previous interview sending a letter to Larry Roberts, when he was already at DARPA, talking about your interest in networking. What happened after that? How did you get involved in working with the ARPANET?

KAHN: Well, I had been doing my own work in that area while at BBN, unaware of the fact that DARPA was interested in networking at the time. So I had a set of memoranda that I had been generating on various aspects of networking. It was for my own research project. I guess the first that we really got involved was Larry responded to that letter by inviting me down to chat and find out who this stranger was that sent him this random letter out of the blue. I learned from him at the time that he was seriously thinking about creating a network across the country. You have to remember, I was a fairly young person just on a leave of absence from MIT at the time, so I guess I hadn't fully digested the idea that one would actually be able to make one of these things really happen.

O'NEILL: Do you remember what time this was, was this early 1967? Roberts came to ARPA in very late 1966, December of 1966.

KAHN: Sometime in 1967. That was my recollection.

O'NEILL: So Roberts told you at your first meeting that he was interested in actually getting a network project started?

KAHN: He had gone to DARPA (then called ARPA) to actually make such a network happen, and he was in the planning stage for it. I thought that some of the ideas I had could be of use in that project. We chatted a bit about what I had been up to. There had really been no very direct interface with him up until the time that DARPA actually issued an RFP. I guess that was what he had been working on at DARPA. DARPA issued an RFP back in... some time in the summer of 1968 - June or July, if I recall. And I remember getting a copy of that just as soon as it came out, and we at BBN put together a proposal that eventually won.

O'NEILL: So between the time that you talked to Roberts and the time that the RFP or RFQ came out in mid 1968, you didn't have any interaction with the community who were working on developing the RFQ? There were a series of meetings that have been referred to and different people being involved coming up with the requirements for the RFQ in detail, in the detail that came out eventually.

KAHN: Roberts certainly had asked me for any of the papers that I had. After we had our first meeting, he said, "You know, if you can share any of your stuff with me, I'd be happy to see it." Well, I think I sent him some of the stuff that I had been working on. But I had no direct interaction with Larry during that period, at least that I can recall. I think he met with people like... there's a fellow named Frank Westervelt from Michigan, there's a guy named Elmer Shapiro from SRI. Larry had even given a contract, I think, to SRI to help him draft the RFP.

O'NEILL: I have seen a reference to some of your papers. So Roberts had your papers and they may have been circulated within that group, but you didn't...

KAHN: He certainly had my papers. In fact, as I was writing stuff I generally would send it to him. I can even recall sending him one paper that basically said whatever he did with the network, he ought to have a long line involved in it up front so he could find out whether the network would work with long distance lines or not. Because everybody was thinking about small scale experimentation. I was afraid if they only did a small scale

experiment that it might work in the small but fail in the large. So I felt they ought to set the initial structuring of it to be consistent with what the eventual network would look like. Worry about the network on the scale that you really want to operate it. You know, if you are going to deal with real errors make sure you are going to encounter approximately the same number you would expect in a reasonable number of configurations. It seemed to me at the time that the longest lines you would ever want in a network would be roughly a few thousand miles long. So put in one or more of them initially, just to see how it would work. Larry was talking to lots of people at that time. I mean I don't know all the people he would talk to, but having been in the ARPA office later, the number of people that you come in contact with could be enormous. He could have been talking to thousands for all I know.

O'NEILL: When you started working on the ARPANET at BBN from 1969 to 1971, how did you see your role in the group that was developing the IMPs and implementing the network?

KAHN: Well, originally I wasn't even planning to get involved in the implementation. You have to remember I started out as a professor in the EE department at MIT. In fact, the main reason I had gone to BBN was to spend some time working on more practical kinds of things than just mathematics. The thing I had chosen to work on was networking. Having now written the networking proposal, I was fully prepared to get back into research and leave it up to other people to go build this thing. But it became very clear to me, shortly after we had gotten the award, the set of issues that were involved was very complicated - relative to what you normally find in a typical engineering project. In fact, my being involved in it was actually far more important than I originally thought because I really played a key role in bringing it all together architecturally. It was a design role that was very essential to what was going on.

O'NEILL: You are talking about during the time of writing the proposal?

KAHN: Yes, and afterwards too. I took the lead in putting together the technical part of the proposal. The whole thing was sent out of the systems division that Frank Heart ran. There were a lot of people worrying about things that I didn't know much about, like subcontracts with the computer vendors and the details of building the hardware because that was not really my specialty. In terms of the conceptual parts of it, that is really where I was main contributor, I think. My notion originally was "Help them get the award; let them go build it and I will go back to doing what I was doing." But I pretty quickly came to realize that that just wasn't a practical notion. There were too many things that had to be thought about. System design was going to be a continuing factor; it wasn't just something you did once up front and then forgot about. There were really a very small set of us that were directly involved in the actual development of the network, and I was one of the key people. Some of the others were Severo Ornstein who was principally responsible for the hardware, Bill Crowther who was principally responsible for the software. He was helped by a fellow named Dave Walden who subsequently became president of BBN Labs. I don't know if you ever talked to Dave.

O'NEILL: Yes, I did.

KAHN: Under Frank was Bill Crowther, under Bill Crowther was Dave. Dave eventually became more senior than Frank in the organization.

O'NEILL: Did you feel any external pressure to continue with this networking stuff because you were so involved in it, or was it your own personal decision?

KAHN: By external pressure what do you mean?

O'NEILL: Well, were the other people that you had worked with on the proposal urging you to continue on with it, or was it just your own decision to...

KAHN: Well, organizations are complex structures. The way BBN was structured at the time, the research division had a different outlook than the engineering division. One was more focused on getting things built on schedule, and the other was more concerned with inventing the future, if you will. That is really where I had spent most of my career, more in the basic research line of things. And having gotten this far in the proposal stage, my idea was to go back into research. That is sort of what was in my psyche at the time. Frank and some of the people in his group had made it very clear that they would like me to keep on with the networking part. They probably had more reason to believe that I would be able to help than maybe I did at the time.

But I think what finally triggered my going over there was an activity that involved simulation. Along with another colleague at BBN, we developed a simulation program that could be used for evaluating network performance. Most of the really interesting issues to me about networks were how they were going to work, because this was a really new area. Nobody had any experience in knowing how networks would actually function in practice: how did routing algorithms work, how did congestion control algorithms work, how did flow control work, how did buffering work. All of those kinds of things were critical issues back then. The simulation system was developed to give us a visual clue. I mean, this was a rather innovative development because - remember this was at time when interactive graphics were not par for the course. I mean, there were no such things as workstations. And time-sharing was pretty new at the time as well. So to be able to get an interactive environment with graphics displays of networks was a major coup in its own right. So we had a very powerful facility to be able to use, and my thought was to use it to explore some of the issues.

BBN at the time was structured into two parts. Leo Beranek ran one part of the company having to do with acoustics and the old line of physical sciences work that BBN had originally started in. They did acoustical design of concert halls and things like that. The other half of the company was run by Jerry Elkind who subsequently left BBN, briefly went to MIT, then went out to Xerox for many years. Jerry had several divisions underneath him, one of which was Frank Heart's division. So he was Frank's boss. In fact, he also had a division of his own that he ran, and a third one, I think in psychology or something like that or man/machine

interaction. And I think it was Jerry's perspective that Frank and his people needed me. So rather than tackle the issue any other way, his approach was to say to me, "I think working on this networking stuff is fine, but you ought to get your support directly from Frank since he has the one contract at BBN specifically to work on networking." The way accounting was handled at BBN in those days, you had to have charge numbers for everything. When I went to Frank to find out what number to charge my time to, he was basically of the opinion that the only way I could charge to the contract was to be in his division. Apparently, he did not want to spend money in other divisions around the company. In any event it became clear to me, very quickly, that the only way I was going to get supported to do networking research was to be involved with that group. I think this was just Jerry's way of saying, "I would like you to do that."

O'NEILL: This was all after the proposal had been accepted?

KAHN: After the proposal was won, right. So in some sense I found myself sort of like somebody on a slide being asked "Why did you go down the slide?" Well, once you are on the slide gravity just sort of takes you. This was a situation where I would have been perfectly happy to do the work in the research division. So I, in fact, picked up my office and moved over to the systems division. Formally, I was still in Jerry's division, only physically in another location. I never knew quite how the payroll was handled at that time, so I can't tell you all the machinations.

O'NEILL: Did you consider yourself a consultant to the network group?

KAHN: No, I was clearly a part of it. Let me put it this way: the most visible public aspect of the whole project was the guidance that went out to everybody as to how to connect to the network and make it work. That was a fairly thick document that we wrote early on. It was called RFC 1822 or BBN report 1822. I wrote that.

KAHN: Yes, and afterwards too. I took the lead in putting together the technical part of the proposal. The whole thing was sent out of the systems division that Frank Heart ran. There were a lot of people worrying about things that I didn't know much about, like subcontracts with the computer vendors and the details of building the hardware because that was not really my specialty. In terms of the conceptual parts of it, that is really where I was main contributor, I think. My notion originally was "Help them get the award; let them go build it and I will go back to doing what I was doing." But I pretty quickly came to realize that that just wasn't a practical notion. There were too many things that had to be thought about. System design was going to be a continuing factor; it wasn't just something you did once up front and then forgot about. There were really a very small set of us that were directly involved in the actual development of the network, and I was one of the key people. Some of the others were Severo Ornstein who was principally responsible for the hardware, Bill Crowther who was principally responsible for the software. He was helped by a fellow named Dave Walden who subsequently became president of BBN Labs. I don't know if you ever talked to Dave.

O'NEILL: Yes, I did.

KAHN: Under Frank was Bill Crowther, under Bill Crowther was Dave. Dave eventually became more senior than Frank in the organization.

O'NEILL: Did you feel any external pressure to continue with this networking stuff because you were so involved in it, or was it your own personal decision?

KAHN: By external pressure what do you mean?

O'NEILL: Well, were the other people that you had worked with on the proposal urging you to continue on with it, or was it just your own decision to...

interaction. And I think it was Jerry's perspective that Frank and his people needed me. So rather than tackle the issue any other way, his approach was to say to me, "I think working on this networking stuff is fine, but you ought to get your support directly from Frank since he has the one contract at BBN specifically to work on networking." The way accounting was handled at BBN in those days, you had to have charge numbers for everything. When I went to Frank to find out what number to charge my time to, he was basically of the opinion that the only way I could charge to the contract was to be in his division. Apparently, he did not want to spend money in other divisions around the company. In any event it became clear to me, very quickly, that the only way I was going to get supported to do networking research was to be involved with that group. I think this was just Jerry's way of saying, "I would like you to do that."

O'NEILL: This was all after the proposal had been accepted?

KAHN: After the proposal was won, right. So in some sense I found myself sort of like somebody on a slide being asked "Why did you go down the slide?" Well, once you are on the slide gravity just sort of takes you. This was a situation where I would have been perfectly happy to do the work in the research division. So I, in fact, picked up my office and moved over to the systems division. Formally, I was still in Jerry's division, only physically in another location. I never knew quite how the payroll was handled at that time, so I can't tell you all the machinations.

O'NEILL: Did you consider yourself a consultant to the network group?

KAHN: No, I was clearly a part of it. Let me put it this way: the most visible public aspect of the whole project was the guidance that went out to everybody as to how to connect to the network and make it work. That was a fairly thick document that we wrote early on. It was called RFC 1822 or BBN report 1822. I wrote that.

O'NEILL: Okay. When you mentioned before about the simulation program that you had written, you mentioned somebody else was involved in that. I did not catch the name.

KAHN: Warren Teitleman. It was just a small project that he and I worked on. It was an interactive graphics simulation system built on top of some other software that he had developed. Warren was absolutely crucial to that.

O'NEILL: When was that done? Were the simulations done as part of preparing the proposal, or was that later?

KAHN: It was just an interesting idea that we had that we explored. Probably, I would guess, in 1967 or early 1968, but not for any proposal.

O'NEILL: While you were working on the network project, what was your interaction like with ARPA and with Larry Roberts?

KAHN: It became very tight. Larry and I developed a very good personal relationship. In fact, he later hired me to go to DARPA. It was a very interesting relationship because Frank Heart, who ran the group in which I was physically residing, also had an interface with Larry. So Larry was, in effect, talking directly to two different people in that group. I was not usually privy to his discussions with Frank, so I don't know what the two of them talked about. But it was a separate, and in some cases, independent channel, because Larry would often call me as easily as he would Frank. I had quite a number of interactions with him, but they were usually on technical matters, planning for the future, "What do you think we ought to do about this," "I got this thing I need to respond, can you handle it?" That sort of stuff. And oftentimes he would send things up for me to think about, you know, go talk to a group of people on such and such, or whatever. We had a pretty good

relationship. Of course, I later became director of the same office, so I had the same position that he did. But that was some number of years later. Larry actually left DARPA in 1973.

Steve Levy and I... Steve was a vice-president of BBN at the time. Steve was involved in new financial opportunities for BBN at the corporate level. I had been pushing BBN to get into the commercial networking business for many years, going back to 1969. Originally it was a recommendation to a small panel that Jerry Elkind or Steve Levy had been chairing. I made subsequent recommendations along the same line to the senior management at BBN, but they had never done anything with them. And, in fact, one of my reasons for accepting the DARPA job was because it just didn't look like BBN was interested in capitalizing on any of the technology we had developed.

Some time in the 1972 time frame, shortly after I had agreed to go to DARPA, one of the people in Frank Heart's group that he had just hired from the Pentagon, a fellow named Lee Talbert, resigned from BBN and took some people with him to set up a small company called Packet Communications, Incorporated - PCI. PCI filed the very first 214 filing with the FCC to offer packet network services. They got approval for that, but their effort subsequently failed for a number of reasons. But right after they left BBN, I recall Steve Levy coming into my office and basically saying that BBN had just revisited the whole networking situation and they thought my idea was now timely, and they were all ready to go make something happen. So over the course of that summer, sometime between June of 1972 and the time that I finally left to go to DARPA - I guess I left in October and got to DARPA in November with a short stop along the way at an ICCC conference - Steve and I worked together to set up Telenet. Steve and I had a lot of interactions on it. The first two employees of Telenet were Stu Matheson and Phil Walker who were two people that I brought in to help with some of the initial planning. I don't believe Frank Heart was involved in any of that activity at all. Steve Levy could tell you exactly what happened if you talk to him. He is now chairman of the board of BBN.

O'NEILL: Do you know why BBN wasn't more receptive to the idea earlier? You said you mentioned in originally in 1969. Were you given reasons at that time?

KAHN: Not really. You have to understand what kind of an organization BBN was. It changed fairly radically somewhere along the line. BBN was a kind of hybrid version of Harvard and MIT in the sense that most of the people there were either faculty or former faculty at either Harvard or MIT. If you've ever spent any time at either of those places, you would know what a unique kind of organization BBN was. A lot of the students at those places spent time at BBN. It was kind of like a super hyped-up version of the union of the two, except that you didn't have to worry about classes and teaching. You could just focus on research. It was sort of the cognac of the research business, very distilled. The culture at BBN at the time was to do interesting things and move on to the next interesting thing. There was more incentive to come up with interesting ideas and explore them than to try to capitalize on them once they had been developed. I think that the administration at BBN would have acted on my proposal earlier if they could have figured out what to do, but the payoff wasn't sufficiently near, and the understanding of what to do with packet networks didn't really exist in industry. So they just sort of let it go. I think "It was ahead of its time" is the right way to put it.

O'NEILL: So until they were prodded in 1972 by having another company starting to capitalize on it...

KAHN: Well, composed of their own people... a non-sanctioned spin off. I think BBN was protecting its interests more than anything else. But they suddenly decided that this was something they really ought to do. I think it was reactive largely. They set up Telenet, which lost money for a long time. I remember, Steve and I talked about Larry as being a potential guy to run it. I just wasn't interested in getting into a business setting at that time. I really wanted to stay with R&D. Somewhere in the mid-1971 time frame, I picked my office up and moved back to the research division again, because we had largely gotten to the point where the networking development effort was stable. We had done the critical phase of what we needed and from then on it looked like whatever else... it's like getting to the moon. Once you get there the first time, you at least know that it

is doable, but there was still a lot of development work that needed to be done to improve the network performance. I wanted to move into other areas, like understanding what to do with the networks - maybe some of the more basic research aspects of it. So I went back into the research group, but that only lasted for about a year. Then, ultimately, I finally went down to DARPA. I showed up there in early November of 1972.

O'NEILL: When you were working out the plan for Telenet with Steve Levy did you recommend Larry Roberts? Did you talk to Larry Roberts about what was going on?

KAHN: I never talked to Larry about it, but I did recommend him. Steve had asked me on several occasions who might be a good choice to run it. It seemed like a very natural thing since Larry was the guy who was the prime mover of the network in the first place. While I was involved in most of the technical parts, Larry was the political mastermind during the project. He got a lot of support from Bob Taylor, who ran the DARPA office when Larry first arrived, and who is now at DEC. I don't know if you have interacted with Bob.

O'NEILL: He has been interviewed for this project.

KAHN: Bob hired Larry to run the network project. Larry was hired in as a special assistant of some sort before he became the office director. Bob Taylor was the one who created the environment for him and gave him the charter and the funding.

O'NEILL: Did you know that Roberts wanted to leave the office at that time? Or were you just thinking of people who would be good for Telenet?

KAHN: I thought about who might be good at that time. He seemed like the right choice. I had no idea that Steve Levy actually went down to talk to him. Steve apparently talked to him sometime in early 1973. Larry

announced in May of 1973 that he would be leaving later that year. He actually departed at the end of September of 1973.

O'NEILL: Okay.

KAHN: So he was actually at DARPA for four or five months after he announced that he was intending to leave. When he left at the end of September, his successor had just been named. It turned out to be J.C.R. Licklider from MIT who came in for a second tour. He started the IPTO office in the first place back in the early 1960s and actually showed up to stay in January of 1974. So there was a period of three or four months where the office really didn't have a permanent director.

O'NEILL: During this time were you running the network project?

KAHN: *The* network project. You have to understand that the way DARPA was organized, it had a lot of different programs - and the ARPANET was one program. One of the things that I chose not to do was to run any of *the* network project. Having just come from BBN, I wanted to get into new things. In fact, the agreement that I had when I went to DARPA was I would set up a program in flexible manufacturing. I was intending to make a clean break from networking. Apparently quite a bit of money had been set aside, or planned, or budgeted, for the new program, but when I got there it had been canceled by the Congress. So the program somehow disappeared in real-time right in front of my eyes. I remember Larry coming to me and saying "Look, I know you didn't want to work on networking anymore, but you know more about it than anybody else around and that's where our main efforts are going to be for the next several years, so why don't you just go do that." So I did, but I got into all the new efforts. The actual running of the existing ARPANET program was left to other people. Of course, people used to talk to me about the network, so it wasn't exactly as if I was isolated from it. However, I was never the official program manager for any of the old existing programs.

O'NEILL: I would like to back up and ask a few more questions about your time at BBN. We talked about your interaction with Roberts and your direct connection with him. What about with the other contractors working on the ARPANET? I am thinking of the people at the Network Measurement Center, and the...

KAHN: I had most of the external interactions with people in the community. You have to remember that here were a bunch of people in the research community who were expecting a network to show up and didn't know what it was. That is why I wrote that document [Report 1822]. I also used to talk to them quite a bit by telephone. Dave Walden and I went out and spent quite a bit of time on the west coast trying to do some of the early debugging of the network. The very first installation was out at UCLA, where Vint Cerf and I had our first opportunity to work together on measurement and testing.

TAPE 1/SIDE 2

KAHN: Although I had quite a number of interactions with the different groups who were involved, they organized themselves pretty well, partly through the workings of the ARPA office itself. One of the critical issues that came up then was "suppose these machines were all connected to the net somehow - how would that all work?" I mean, how is a machine going to be able to move a file from here to there? How would that work at the host level? And a mechanism got set up to address this issue. I think Barry Wessler, who worked for Larry Roberts, was one of the key guys that put it in place, along with a fellow named Steve Crocker who was then at UCLA. Steve organized something called the Network Working Group, whose job was to get host-level people together, to discuss what they were going to do, what things they had in common, and what standards were needed. In particular, they were to come up with a protocol (or set of protocols) that would allow machines to talk with one another. Steve ended up chairing and coordinating that whole process. But it was, in fact, a very broad group activity. And it started a series of working notes called RFCs which still continue to this day.

O'NEILL: Were you a part of the Network Working Group?

KAHN: I sat in on some of the meetings, but I did not consider myself a prime mover in that activity at all. Often times I was there to represent the network, to tell them what we were doing, or explain things that were written in some documents. Occasionally I would contribute a note to the series. But this was mainly a group of people who were more involved at the host level, and I was more involved in building the net that would move the bits between the hosts. Vint Cerf was definitely involved at that level. In fact the original paper that was written about host protocols was co-authored by three people: Steve Crocker, Vint Cerf, and another fellow named Steve Carr who was at Utah. Carr later just disappeared from the scene and I never heard of him again.

O'NEILL: You mentioned that you, out of the group of the people at BBN, were doing a lot of the external interface to the various other contractors. What about other people interested in networking?

KAHN: Frank Heart had a lot of interactions with DARPA. And we occasionally would have meetings where we would bring together various people. Like we hosted some of the Network Working Group meetings, and there were a lot of people from BBN that would sit in on those meetings. So a lot of people from the group had a lot of meetings with a lot of those people. I'm not trying to say that I was the only one. But if you look at a typical week, I mean, I would be on the phone typically every day with many of those groups talking about what was happening, what is the schedule for this, or they would be calling about such and such or writing the interface document for them to use. I had a responsibility to do that, but a lot of people were also involved from time to time.

O'NEILL: That was leading me into a question about others interested in networking, and in particular I had in mind Donald Davies and his group at NPL in England. They traded visits with BBN, as I understand it. Why don't you describe for me what the interaction with them was like.

KAHN: Well, I had met with Donald a number of times. I think it was largely in the early 1970s. Donald had gotten the idea of packet switching into his head in the 1960s, I don't know from where. He set up a little project at the National Physical Laboratory in Teddington, just outside of London where he worked, to explore it. He acquired a Honeywell machine and outfitted it with some special purpose interfaces so he could connect a whole bunch of terminals. He actually built a one node network. You could type text on one terminal and designate which other terminal you wanted it to come out on, and the text would print out on that other terminal. That was the extent of his network, I believe.

O'NEILL: Would that be a similar idea to a terminal IMP, where you are connected just via your terminal versus a host computer?

KAHN: Except that he never really had anything external to that one machine. Yes, you could think of it like a terminal IMP. In fact, that's very much what it was except that the terminal IMPs had all these network protocols in them that he probably didn't have in his implementation because it was fairly simple early implementation. Donald Davies was a very creative guy; he thought a lot about interesting ideas of how networks should be built. He clearly had the concept in his head of what packet networks ought to look like, and he had done it independently in England. I believe Larry Roberts will probably tell you that Donald had a big influence on him. In fact, I think the term packet switching came from Donald's work. We had been calling it message switching before, on the notion that what computers did was send messages into these network nodes, which broke them up into packets, sent the packets through the net, and put them back into messages at the destination. So we were thinking of the network as a message switching system - but the term packet switching eventually stuck.

O'NEILL: Do you recall when it started to be called that, when the terminology changed?

KAHN: Well, we wrote the very first paper on the ARPANET IMPs in 1970, and we did not use the term packet switching in that paper. We also called it the ARPA Computer Network, so the term ARPANET was not yet in use, and packet switching was not quite in the vocabulary yet. Not that the term packets were not. I mean we referred to packets, which would be switched from here to there, but the term packet switching was a term we didn't actually use. I think that it was probably a year or two later that it actually stuck, maybe 1971, or early 1972.

O'NEILL: I've been trying to identify when it actually appeared. I haven't been able to find anything.

KAHN: Leonard Kleinrock, Howard Frank, and I wrote a paper on "Lessons Learned," that was 1972. That was the first attempt to bridge the gap among theory, simulation, and engineering. We came at it from three different points of view. Len specialized in theoretical analysis. Howard ran the group that was doing all the topological analysis; that was all done by simulation. Where should one put the lines and how much bandwidth was needed - things like that. I was the focal point for the engineering aspects of the network design. When we all got together to try and figure out how to write what we had learned, we all would agree, "We don't know how to build nets today that are bigger than about 30 or 40 nodes." And everybody would say, "Right." "Can we craft a sentence that says why?" And I would say, "The reason we can't is because the lines fill up because you are passing too much routing information back and forth." Howard would say, "The reason is because today's simulation systems can't handle more than 30 or 40 nodes. You can't get answers back." Len would say, "The reason is that the denominator of this formula goes to zero when rho equals such and such." Everybody had a different view of what was important, because they were focusing on different things. That was actually a very interesting paper that we wrote. It was the first attempt to get the lessons learned out in writing. It would be interesting to look in that paper and see if the term packet switching shows up.

O'NEILL: I don't believe it does.

KAHN: You have read that paper?

O'NEILL: Yes, I have. I read it before I was looking for a change in terminology. So it could be, but I don't think it was.

KAHN: It is possible that it wasn't around yet even by 1972, although I would doubt it. It may be that we didn't have any reason to use the term. I don't think we used the term message switching then.

O'NEILL: I will have to look again.

KAHN: I am sure that had long gone by then. So some time in that time frame packet switching terminology just sort of slipped in as did the term ARPANET. It became part of the culture. It's like Internetting. When Vint and I wrote the original paper on internetting, I don't thing we actually used the term internet, which became as much a part of the jargon as ARPANET did. We wrote about internetworking, we talked about internetting. I mean the sequence of letters showed up, but not explicitly as "Internet." That was, again, something that just sort of came upon us during the 1970s.

O'NEILL: Getting back to Davies and NPL, were they the only other group that was interested in networking around this time? Were you discussing networking issues with people in Europe or people at other places who weren't host sites, who weren't directly involved in the ARPANET?

KAHN: Well, there were a lot of people who were interested in the topic. In some sense, it all started with the work that Paul Baran did at Rand. Paul used to work for Keith Uncapher, who now works for me. Paul was and is a very brilliant fellow, still as agile today, mentally, as he was back then. But Paul was not a technologist at the time in terms of actually trying to build something. He had a concept for how a network might work in a hostile environment. His concept was based on breaking messages into little blocks, each with an address for

transmission. He originally called it Message Addressed Block Switching, or something like that, back at the time.

O'NEILL: Adaptive Message Block Switching.

KAHN: Adaptive Message Block Switching is an awkward composite name. Paul was thinking of that as a survivability mechanism for communications in times of nuclear war. You have to remember the context in the early 1960s. Nuclear confrontation planning was consuming the military. My cousin [Herman Kahn] wrote two very famous books called *On Thermonuclear War*, and *Thinking about the Unthinkable*. How do you plan constructively in the thermonuclear age?

O'NEILL: This is the early 1960s?

KAHN: This was Herman Kahn's work in the early 1960s. And Paul did his work in that context. His idea was if you ever had a network that could break messages up into addressed blocks and send them independently, then if a block could find an available path, somehow it would probably get to the destination. Whereas if only a specific route could be used, maybe that route would be blown up, or otherwise inoperable. I think Paul was focusing mainly on survivable voice communications, which was a big Air Force problem at the time. Not too much on computers, although he did understand you could link computers that way. He had this notion of low cost discardable electronics to do the switching. He hadn't quite focused on exactly how to do this, however. You have to remember that VSLI didn't exist in those days, and so there was probably no feasible implementation. Sort of like Jules Verne imagining the submarine in the 1800s. The only difference was Paul's conception was within less than a decade of actual implementation. And it became economical when the mini-computer was invented. I believe he had no detailed notion of how it might really be implemented, except somehow electronics might one day be cheap enough to do the task. He did not develop the notion of protocols,

which really became a major aspect of networking. He didn't seem to be thinking about the computer and the computer communication problem as much as he was voice communications.

O'NEILL: Was he still in the community during the time? I had the impression that his work in networking had pretty much been completed and he moved on to other things by the late 1960s.

KAHN: Paul was not an active player in the network development phase. He is credited with a major contribution, and it is probably fair because he had a lot of new and innovative ideas. But he was not part of the process that made it actually happen, unless you somehow view his early series of reports as being part of the process. In fact, most of us were not even aware of his work until we were significantly into the process. I mean I didn't even know about Paul's series until probably some time into the mid-1968 time frame.

O'NEILL: How did you become familiar with it? Why did it show up in 1968?

KAHN: Probably Larry had mentioned that there was a report series that Rand had generated. The Rand work was not published widely because it was mostly for defense purposes. Paul's reports were first published internally at Rand. So it could very well have been hard for people to get hold of. I don't think the report series was published externally. If they were, I don't know where. However, aspects of this work were published in the IEEE transactions, if I recall.

O'NEILL: The Rand memos did get sent out publicly. They were more than just internal memos.

KAHN: Did they? Okay, I never saw them. Larry says he wasn't even aware of Donald's work in London when he began working on networking. I can believe that. I mean, Larry was just... You have a bee in your bonnet and you go off and do it, and you find out "Oh, that guy was also doing it. Let's find out what he was up to."

you use and what information gets passed back and forth. It is one thing when you plug a socket into the wall and electrons flow; it is another thing when you have to figure out for every electron which direction it takes, and how many of them, how many per unit time; what happens if that machine is down, and these buffers are full. So it was a rather different kind of question. It wasn't quite theoretical, it was more of a system design issue, an architectural issue to a protocol design, in part. That whole field of protocols and architectures was really in its very early infancy back then. Most people didn't understand that it was a discipline of sorts at that time.

O'NEILL: Okay.

KAHN: That's really what the focus was - things like how do you do routing, what do you do when congestion occurs. That was one of the main points of contention early in the days of the network. It was my contention that we had to worry about congestion and deadlocks. What do you do when the network just fills up? Things might come to a grinding halt. I was busy at work designing mechanisms that would prevent that from happening or, if it did happen, that would get you out of it. And the prevailing feeling of my colleagues was it's like molecules in a room; don't worry about the fact that you won't be able to breathe because the molecules will end up in a corner. Somehow there will be enough statistical randomness that packets will just keep flowing. It won't be able to somehow block itself up. It was a significant point of contention. When Dave Walden and I went out to the west coast at the at the end of 1969 or early 1970 to test the net the very first thing that we did was run deadlock tests. And the network locked up in twelve packets. I had devised these tests to prove that the network could deadlock. There was no way to convince anybody else, particularly the person writing the software, that the network was going to deadlock - except by doing it. And even that wasn't sufficient as it turned out.

O'NEILL: The experiment wasn't sufficient to show that it would happen?

KAHN: Yes. Because everyone thought the experiment must have been contrived. They still didn't think it could happen in actual use. It was about nine months later when some internal work that Bill Crowther of BBN had been doing on simulation led him to the same conclusion and then it all changed. But it was rather contentious for a while.

O'NEILL: Okay. Let's move on to 1972 and the ICCC demonstration of the ARPANET. How did that idea come about? I know you had a large role in organizing it. Was it suggested to you as a project? How did it get started?

KAHN: Well, I can give you my view of it. I had been thinking for a while that it would be useful to have some kind of a demonstration of the net. My recollection is that in early 1971 I had suggested to Larry Roberts that we try to put on a demonstration at the Spring Joint Computer Conference in 1972. That never developed, but Steve Crocker called me back soon thereafter and said there was an opportunity to actually put on a demonstration at a new conference called ICCC in the fall of 1972. Larry had made a deal with the conference organizers that if they were willing to put up the space he would somehow arrange a network demonstration. I mean, I think I proposed to Larry that they do a demonstration. I think the fact that we did it at ICCC was his idea. It is entirely possible he thought it up independently. It doesn't really matter. The fact is it was clear to a lot of people that having a network demonstration was a good idea. Because in the middle of 1971, despite the fact that the ARPANET had been able to deliver packets for a year and a half at that point, there were almost no useful interactions that were taking place on the network.

O'NEILL: Why was that?

KAHN: Well, it would be like having a highway system that was perfectly capable of handling automobile traffic, except there were no cars around, or no on-ramps and off-ramps. The reality was that the machines that were connected to the net couldn't use it. I mean, you could move packets from one end to the other. You could

run all the test programs you want. The network nodes could even send test traffic. I could sit at the teletype at one place, connect it to an IMP and talk to somebody at the other end, but none of the host machines that were plugged in were yet configured to actually use the net. That was one of the main reasons why I was thinking that a demonstration would be useful. It would put some pressure on the community to make the connections useful.

We had a meeting at MIT during the summer of 1971 at which I recall putting a matrix up on the board. The rows were the names of host machines. Like number one was UCLA, number two was SRI, number three was UCSB, number four was Utah, and so forth. So we listed all the machines, and at that time there were probably something on the order of twenty-ish machines available. The columns indicated the minimal things you could do over the net. Like you could connect, or you could actually get a herald back - herald meant you could get the protocol to open a connection to a process. Or that you could actually run a program, by "a program" I mean at least one. You could do something once the herald came up. Or perhaps you could run an arbitrary program. We had about ten columns that denoted various stages of ability to do something. It would be the equivalent of, "You could drive your car up to the front door. You could actually walk from your car and grab hold of the door handle of the front door. Or, you could actually open the front door and walk in the lobby. Or, you could walk from the lobby into the office, or you could actually sit down in an office." It was sort of that level of penetration into the machine. And we filled out this whole matrix and there were perhaps six entries in the first column, and maybe three entries in the next column, and maybe one entry in the third column, and nothing else. So the fact of the matter is that - maybe it was ten, five, and two, I don't remember. But the fact of the matter is that this was an incredibly sparse matrix. The hosts basically couldn't use the net. My thought was "Look, we've got to fill in this table, and make the network more usable; otherwise the demonstration is going to be a bust." That was my logic. But at the same time a main reason for having the demonstration in the first place was to make it all work. So that was the process that actually got the utility of the net up from ground zero to some reasonable level.

O'NEILL: What, in your opinion, was responsible for the sparseness of the matrix?

KAHN: Well, you have to remember every site needed to do something to get connected. One, they had to figure out how to build a hardware interface between their computer and the network. Almost all that hardware was special purpose. Sometimes that took six months or a year; in some cases they built it in-house. Then behind that you needed to build all the software. Somebody needed to implement the network protocols in those machines. That was fairly complicated stuff. I mean it wasn't like you bought a package from Lotus, back in that time. It was probably some creative scientist on the staff that wrote the code, which had to be interfaced to the operating system. This was uncharted territory, absolutely uncharted territory. And people needed some motivation to get it done. The demonstration was actually motivation to do it.

O'NEILL: Okay, so the demonstration really worked in two ways. Before the demonstration it worked to give more sites incentive to put cars on the highway.

KAHN: It gave them an incentive because it set a target to shoot at with a public deadline, a target that really couldn't be put off.

O'NEILL: Then the actual demonstration was to a wider community who could look at the ARPANET and see that it was functioning.

KAHN: That's right. That was the event that made the world take notice of packet switching, because up until that point you couldn't see it anywhere. All you could do was read an arbitrary abstract paper somewhere that said, "Here is this new way to do computer communications." A lot of people were skeptical in the early days. I mean, breaking messages into packets, reassembling them at the end, relying on a mysterious set of algorithms, routing algorithms, to deliver packets. I'm sure there were people who distrusted airplanes in the early days. "How are you going to ensure that they are going to stay up?" Perhaps this was the same kind of thing. But

ICCC was the event that made... I think this was the watershed event that made people suddenly realize that packet switching was a real technology. We had thousands of people who went through that - I don't know what the exact number is - at least a thousand who went through that particular exhibit. Which reminds me, I had some audio tapes from that exhibit that I was going to try to get hold of.

O'NEILL: Yes, they had been mentioned in the previous interview. That was one of the things I was going to ask you, off-line, if you have any of those available.

KAHN: I'm sure I have it, I just need to find it.

O'NEILL: Okay. After the 1972 demo you went on to work at the DARPA office as program manager, and you explained how you had not gone to work on networking but ended up working on networking projects. How did the idea of the Internet come about? What was the problem that you were trying to solve?

KAHN: One of the very first efforts that I undertook when I got there was to develop a technology that was like the ARPANET technology except it was based on radio communication. So the idea was to build a net that could link a number of entities, computers let's say, where each of these entities could be, in principle, on a moving platform although they could also be stationary. But the communication was by radio. So instead of machines linked to each other by wires, they would be linked by broadcast radio. A given node would broadcast its radio packet and a nearby node would pick it up and relay it. There had been an effort supported at the University of Hawaii in the late 1960s, early 1970s called the Aloha System. I think it originally got its money from AFOSR, then DARPA became a big supporter of it.

TAPE 2/SIDE 1

O'NEILL: You were talking about the Aloha system, that it had been funded...

KAHN: Funded by AFOSR first, then by DARPA through AFOSR, I think. I don't know if they had any more continuing money, but DARPA became the main supporter. They demonstrated for the first time that you could do packet communication by radio. Not that I think, in retrospect, there should have been any question about it, but you've got to demonstrate it for the first time. Aloha was a centralized system. That is, packets were sent from a user's location to the central computer location on one frequency and answers were sent back on another. So Aloha was a one-hop system. The packet radio notion expanded that by taking - it was sort of like extending Donald Davies' experiment with a single node multiplexer. It turned the Aloha idea into a network concept by essentially finding ways to route packets from node to node, to allow nodes to be in motion relative to each other, to deal with the vagaries of an environment where you didn't have a tall central antenna that everybody could see, where people could be in tunnels part of the time, and all the vagaries of trying to build a net like that. I was principally involved in architecting the packet radio concept out of the DARPA office. I mean I undertook a fairly similar effort to the technology effort that I did for packet switching in the ARPANET. The packet radio program had a lot of contractors; perhaps as many as seventeen contractors were working on this project.

O'NEILL: Had this been a proposal that had come in, or this was something that you had decided was important to start? How did it begin?

KAHN: I actually think the idea for a packet radio network was just like the idea for the ARPANET. The ARPANET idea had been floating around the ARPA office before Larry got there. The idea for a radio-based version of ARPANET, and a satellite-based version, had also been floating around the ARPA office for a while. I think there may even have been a small study contract given to one or two organizations to see if it was feasible, before I got to DARPA. When I got there there was money budgeted for a packet radio program, and I undertook to make it happen. The skids were all greased for that. Part way through the first year of the

program it became clear to me that we were going to have to have a plan for getting computer resources on the net. In 1973, mainframe computers were multi-million dollar machines that required air-conditioned computer centers. You weren't going to connect them to a mobile, portable packet radio unit and carry it around.

So my first question was "How am I going to link this packet radio system to any computational resources of interest?" Well, my answer was, "Let's link it to the ARPANET." Except that these were two radically different networks in many ways. I mean, all the details were different. I don't mean conceptually they were different. They were sort of the same genre. Just like, say, Chinese and Americans are of the same genre except one speaks Chinese and one speaks English, one lives on one side of the world, one lives on the other side, they go to sleep during your daytime, etc. The details of the two networks were rather different. The ARPANET ran at 50 kilobits per second and the packet radio system ran at 100 or 400 kilobits per second. One had thousand bit uncoded packets; the other had two thousand bit packets which could be coded. The ARPANET assumed that once you sent something it was delivered with a hundred percent reliability. If it didn't get through the system was broken. The other assumed that much of the time you would never get anything through even though the system was working. The protocols that were designed for the ARPANET wouldn't work over the packet radio net because when a packet entered the packet radio net, the only thing the ARPANET would have told it was where it came from but not where it was going. So the packet radio net had no further information to know where to route it. If a packet got lost along the way, the ARPANET hosts would come to a halt. Well, in a radio net you can get interference and so some loss is natural. So we really had to rethink literally the whole issue of host transport protocols. Vint Cerf and I jointly came up with the TCP/IP concept as a new transport mechanism as part of an architecture for internetworking. DARPA then gave a contract to Vint at Stanford to actually implement the TCP/IP concept - along with small efforts at BBN and at University College London. Vint had the lead for developing the specification.

O'NEILL: What were the other two contracts for?

KAHN: The one at University College London was part of an effort over there to explore remote use of the ARPANET and to implement a TCP/IP protocol on their machine. I think it was a PDP9 at the time. BBN was under contract to build a piece of the packet radio system (called a station) that needed portions of this protocol. That was the first embodiment of the notion of a gateway. We needed to implement a protocol that would work across the gateway. Eventually we all took the Internet technology pieces and created a separate program in DARPA for it. But originally, all that work was done as part of the packet radio program. I subsequently hired Vint Cerf to come to DARPA and actually run the Internetting program. That was in 1976, and by that time we were already three years into it. So he took it from a fledgling effort and turned it into a major national activity.

O'NEILL: Did you go to Vint Cerf with these problems and discuss them with him because he had previously done the host-to-host protocols? Why Vint Cerf?

KAHN: Two reasons. One, Vint and I had worked very effectively on the original IMP when he was at UCLA. So he knew about the ARPANET; he was one of the three people who were involved in actually writing the original paper on the NCP, so he knew that pretty well. I had a good view of what the linkage between the packet radio system and the ARPANET should look like. And I knew what the protocol roughly wanted to be, but there were all kinds of details that needed to be worked out, not the least of which were what are all the issues that you have to worry about in linking this to the operating systems, which those folks had already solved. So I needed somebody like Vint to be able to represent that part. Of course, Vint was very facile at thinking through the rest of it as well. So it just took one session before the two of us were on the same wavelength as to what we needed to do. And he and I just jointly worked it out from there.

O'NEILL: Had you discussed it with other people in the community as well?

KAHN: Not really. Just the two of us. On occasion, we bounced ideas off of other people, but it was pretty much just Vint and myself working in a little room at what was then the Cabana Hyatt in Palo Alto. He had also come to Washington a few times. We met at Dulles Airport and did some work there, and I had gone out to Palo Alto. So we had actually had a few sessions together and then at one point during the summer of 1973 we just sat ourselves down for an entire weekend and just wrote the paper.

O'NEILL: I am interested in your interaction with the military. How involved did you get, while at ARPA, in trying to either interest the military in your projects or meet their needs? What kind of demands were they making on you, -what was that interaction like?

KAHN: You have to remember that the military largely depends on private industry to provide it with technology, and so a lot of what we did in our office at DARPA ended up affecting the military by virtue of its transition through industry. Time-sharing was an example. That eventually became part of something called WMCCS - Worldwide Military Command and Control System. It wasn't like somebody from DARPA went to the military and said "We have time-sharing, do you want it?" and then somebody in the military said, "Hey, great! We'll take it."

[INTERRUPTION]

O'NEILL: We were talking about the interaction with the military and how they started using time-sharing, not directly but by having it available from industry.

KAHN: That was the typical method that was used before by the particular office that I was involved with. This approach was somewhat different from some of the other offices, however. If you are working on, let's say, technology for an improved airplane, the Air Force would probably understand instantaneously not only why they would need it, where it would fit in, but they might be part of the process of helping to develop it. It is harder

to see how some of these advanced computer technologies could have been handled the same way - at least in that timeframe. Maybe now it is different because there are probably more people in the military who have backgrounds and expertise in computer technology. There were very few in the early 1970s. One of the things that happened as a result of the ARPANET experience was we managed to get DCA to take over the day-to-day operations of the net. So here was a case where we actually took a finished piece of technology and literally got an operational part of the Defense Department to run it. So that was a way to actually transfer technology and to make them aware of it. But it really came after the fact. I mean DCA was not part of the actual creation of the technology, although it became probably the largest user in defense of that technology after the fact.

O'NEILL: Were you involved in the switch to DCA?

KAHN: I was one of the key people who helped to negotiate that with them. The person who was the chief scientist of DCA at the time was Harry VanTrees; and Harry was a principal contact at DCA along with Estil Hoversten, who worked very closely with him. In the case of packet radio technology, packet satellite technology, and some others, it was less clear, because there wasn't an immediate need in the military for them. It wasn't clear how, if you developed a technology, it would ever transition into the military. So in the mid-1970s we came up with the notion of setting up testbed programs with the military, with the purpose of taking advanced technology that was in development, continuing to develop it jointly with a military partner, and actually field testing the equipment with them in some settings. We set up a packet radio testbed program at Fort Bragg in North Carolina with the 18th Airborne Corps, I believe. We actually deployed packet radio technology, and they used it in their exercises. So they could actually see what the practical utility was of this technology in real applications.

Now we never got the technology to the point where they could put it on a plane and use it operationally in, say, Germany or Grenada - this is the rapid deployment force - but, in fact, they could clearly see the utility by using it on the base, and they used it for many purposes. One of the most interesting applications was using it to

facilitate loading of the aircraft for deployment. The time it took to develop an aircraft load plan was typically days to weeks. In such a plan, you assume a given number of C5 aircraft, and a number of 130s, and 141s. And each one of those aircraft have special characteristics. So what do you do with your detailed plan if the C5 doesn't show up. You have to figure out how to redo the whole plan with a different set of aircraft. So instead of getting five 130s, you've got seven of them. It requires a different loading plan. The question is how do you rapidly do this replanning with no notice. You might not even know until you are out on the runway to see what planes actually show up. So they were able to use the packet radio technology on the runway to obtain remote access to a computer based loading program. Also to deal with many other interesting problems of which probably the one that was most interesting to me was the issue of how to deal with the constant need to relocate the pieces of a command headquarters in the field. When the headquarters is being set up you probably don't want to drag along tons of cable. An alternative is to put up an antenna and suddenly be in instant communication among all the pieces, which may be hundreds of yards or even miles away from each other.

O'NEILL: Once they started understanding the technology and what it could do for them, did they request changes to it? Was there any feedback? Did their requirements affect how you then proceeded?

KAHN: Yes, in many ways. Dave Russell, who became IPTO director after Licklider left, was very instrumental in making those testbed programs happen. He handled a lot of the interfacing with the folks at Fort Bragg. We set up the second program with the Air Force Strategic Air Command out at Offutt Air Force Base on a similar kind of problem, except in an Air Force context. We had a number of other examples of testbeds. Dave Russell was an Army colonel who understood the implications of all this technology for the military. He was the office director for the period from August 1975 through August 1979. He was a marvelous point of liaison with the military in that he articulated both the technology point of view as well as the military needs. In fact, during that whole period, I think we opened up a greater dialogue with the military in terms of these testbed programs than we ever had just through technology alone. The agency director at the time, George Heilmeier, was very supportive of these testbed programs. So that was a very important development in the history of the office.

O'NEILL: Okay. Let's move on to the conversion effort to TCP/IP, which as I understand it, happened in January of 1983. First of all, were you active in working to get TCP/IP established as a DOD standard in 1980?

KAHN: The answer is yes, except that I did it in the role of Office Director. Vint Cerf was probably more involved than I was, because it was part of his program. But, in fact, a lot of the drive to get the protocols standardized came from the fact that the ARPANET itself was growing. I mean the ARPANET, after it transferred to DCA, suddenly became a vehicle where a lot of other military sites could now get on easily - just go to DCA. It was their normal way - they didn't have to deal with a research agency.

So by the 1980 timeframe there were already a large number of military sites on the ARPANET. And by 1983 the number had grown so large that the need to split it was becoming urgent for many reasons. Many of the military sites were beginning to depend on the net for operational purposes. Their day-to-day business really depended on it, and they felt, I think rightly so, that there were good reasons to keep the military application of it separate from the research and educational aspects. If you have nodes in the network on university campuses and if there were some major outbreak, it would be hard to guarantee that something couldn't happen to the network. I don't know all their motivations, but it was clear already by 1980 that there was real interest in somehow making sure that this net was a stable, dependable facility for the military. And the protocol that was most widely used at that point that could lead to an eventual partition of the network was TCP/IP. There was still quite a dominant use of the old NCP technology. But if there was ever a hope that we could split the net into multiple pieces, we needed to switch over to the internet protocol because connections between multiple nets needed an internet protocol. So the protocol became a DoD standard in 1980 and Vint played a role in that. The sweep of events at the time was such that DoD really had to decide what guidance to give people who were connecting their computers to the net as newer sites came in. "What do we tell them?" So they finally decided to standardize the protocols, because it was really the only game in town at that point in time.

O'NEILL: Do you recall if there was much resistance to adopting a standard?

KAHN: Well, there had been a move afoot in Europe to develop a competing - I say competing although I don't know that officially that was the way it was viewed - set of protocols within the oversight of the International Standards Organization. There were two standardization bodies over there, CCITT that tended to deal with the lower levels, the physical connections, and then the ISO that dealt with the higher levels. I just think there was a sentiment on their part to.do their own thing. They just didn't want to adopt an American development. So they proceeded to develop their own. They came up with a reference architecture in the mid-1970s, which was actually a very important contribution because it gave people a way to think about protocol layers. It had certainly been in our consciousness, but we had never articulated it quite that way, and they did. That was a real contribution.

But their seven-layer reference model didn't say how any of the layers should be implemented. Interestingly enough, in their original seven-layer model, they did not have an internetworking layer because they assumed end-to-end circuits. It was the old telegraphy model of the world, or telephony model. But later on they decided to add internetworking and so their model became more compatible with what was going on in the U.S. They developed their own standards. X.25 was one of the early protocols for interfacing computers to networks. That protocol came out of some of the early needs of the commercial world to define interfacing standards. Four PTTs were involved in originally agreeing on the X.25 standard. There was some pressure by the late 1970s or early 1980s to have us use the ISO protocols, except that they weren't really well defined, and there were few compatible implementations. We were dealing with protocols in the U.S. that had been around for six, seven years, and had been through three or four major versions. A protocol conversion, even if it were well defined, would have cost many millions, if not tens of millions of dollars given all the sites involved. So there was no motivation or incentive to do it. That was really the only competing pressure at the time.

O'NEILL: In terms of organizing the conversion - as I understand it you were quite involved in that. That was a project that you worked on.

KAHN: Well, Vint Cerf had been running the Internet program and he left in October of 1982, just before the conversion. But I picked the program up after he left and actually managed the transition. I handed the project over to Barry Leiner in September of 1983, so actually I managed the Internet program right through the transition for just about a year.

O'NEILL: What were the big problems with the conversion?

KAHN: Well, I think the biggest problem was just getting people to believe that it was real. It is like any major change; it is not real until it happens. We sent messages to everybody, alerting them to the timing and yet one week before we were still getting messages, "Is this really going to happen next week?" or "Let us know if you decide to really go ahead with this." The day after we did the transition, people were saying "Hey, how come I didn't know about this?" or "It was impossible for me to convert; I need another six months." We would say, "You had two years' lead notice; why is it that suddenly you need six more months?"

So we actually had to deal with the fact that not everybody was able to convert on day one. In fact the way we actually handled it was to allow an overlap of both the old protocols and the new ones for a short period of time, and then by special exception we allowed a few sites to continue using the old protocols, particularly where there was no interest on their part in ever converting. Like military sites, where the only party they wanted to communicate with was one other party or for some reason they were passing classified traffic, and there was only one type of device in the world that would do it and it required the old protocol. Something like that. That was by far the biggest problem - just getting people to take it seriously. But managing it was traumatic for a while. I mean, the phone was ringing off the hook every few minutes. Every day someone new would complain, "I used to be able to do this, and now I can't." Shaking it all down was also a problem. Even the places that thought they were going to convert properly suddenly found that while theirs worked with the three or four places that they thought it would, or had tried it out with, it didn't work with some others. So there was quite a bit of time required to smooth out the rough edges.

O'NEILL: It sounds as if you actually used to network to communicate with people. Is that how you sent messages and tried to coordinate? Did you ever try to bring people together?

KAHN: More often than not, critical problems were handled by telephone. But when you are dealing with a lot of people, it is just a matter of keeping them appraised as to what was happening, sending out notices, schedules, debugging problems, coming up with alternatives. We would advise, "Look, not to worry because in two weeks the following will happen, and you'll be able to do this" - helping everyone separately think through their unique situation.

O'NEILL: Who else was working with you on this?

KAHN: At that time? The whole community.

O'NEILL: I'm sorry, I meant at the DARPA office.

KAHN: Just me pretty much during the transition. You have to remember that that office was really small. We had about ten technical people and a few secretaries.

O'NEILL: I thought with something as large as the conversion, you might have drafted a few other people to help out.

KAHN: Not within the office. Before Vint left he had put in place a number of mechanisms to help with the management of the whole Internet process. He had created something called the Internet Coordination Control Board or some such thing.

O'NEILL: Is that the Internet Advisory Board?

O'NEILL: You were talking about the Aloha system, that it had been funded...

KAHN: Funded by AFOSR first, then by DARPA through AFOSR, I think. I don't know if they had any more continuing money, but DARPA became the main supporter. They demonstrated for the first time that you could do packet communication by radio. Not that I think, in retrospect, there should have been any question about it, but you've got to demonstrate it for the first time. Aloha was a centralized system. That is, packets were sent from a user's location to the central computer location on one frequency and answers were sent back on another. So Aloha was a one-hop system. The packet radio notion expanded that by taking - it was sort of like extending Donald Davies' experiment with a single node multiplexer. It turned the Aloha idea into a network concept by essentially finding ways to route packets from node to node, to allow nodes to be in motion relative to each other, to deal with the vagaries of an environment where you didn't have a tall central antenna that everybody could see, where people could be in tunnels part of the time, and all the vagaries of trying to build a net like that. I was principally involved in architecting the packet radio concept out of the DARPA office. I mean I undertook a fairly similar effort to the technology effort that I did for packet switching in the ARPANET. The packet radio program had a lot of contractors; perhaps as many as seventeen contractors were working on this project.

O'NEILL: Had this been a proposal that had come in, or this was something that you had decided was important to start? How did it begin?

KAHN: I actually think the idea for a packet radio network was just like the idea for the ARPANET. The ARPANET idea had been floating around the ARPA office before Larry got there. The idea for a radio-based version of ARPANET, and a satellite-based version, had also been floating around the ARPA office for a while. I think there may even have been a small study contract given to one or two organizations to see if it was feasible, before I got to DARPA. When I got there there was money budgeted for a packet radio program, and I undertook to make it happen. The skids were all greased for that. Part way through the first year of the

program it became clear to me that we were going to have to have a plan for getting computer resources on the net. In 1973, mainframe computers were multi-million dollar machines that required air-conditioned computer centers. You weren't going to connect them to a mobile, portable packet radio unit and carry it around.

So my first question was "How am I going to link this packet radio system to any computational resources of interest?" Well, my answer was, "Let's link it to the ARPANET." Except that these were two radically different networks in many ways. I mean, all the details were different. I don't mean conceptually they were different. They were sort of the same genre. Just like, say, Chinese and Americans are of the same genre except one speaks Chinese and one speaks English, one lives on one side of the world, one lives on the other side, they go to sleep during your daytime, etc. The details of the two networks were rather different. The ARPANET ran at 50 kilobits per second and the packet radio system ran at 100 or 400 kilobits per second. One had thousand bit uncoded packets; the other had two thousand bit packets which could be coded. The ARPANET assumed that once you sent something it was delivered with a hundred percent reliability. If it didn't get through the system was broken. The other assumed that much of the time you would never get anything through even though the system was working. The protocols that were designed for the ARPANET wouldn't work over the packet radio net because when a packet entered the packet radio net, the only thing the ARPANET would have told it was where it came from but not where it was going. So the packet radio net had no further information to know where to route it. If a packet got lost along the way, the ARPANET hosts would come to a halt. Well, in a radio net you can get interference and so some loss is natural. So we really had to rethink literally the whole issue of host transport protocols. Vint Cerf and I jointly came up with the TCP/IP concept as a new transport mechanism as part of an architecture for internetworking. DARPA then gave a contract to Vint at Stanford to actually implement the TCP/IP concept - along with small efforts at BBN and at University College London. Vint had the lead for developing the specification.

O'NEILL: What were the other two contracts for?

program it became clear to me that we were going to have to have a plan for getting computer resources on the net. In 1973, mainframe computers were multi-million dollar machines that required air-conditioned computer centers. You weren't going to connect them to a mobile, portable packet radio unit and carry it around.

So my first question was "How am I going to link this packet radio system to any computational resources of interest?" Well, my answer was, "Let's link it to the ARPANET." Except that these were two radically different networks in many ways. I mean, all the details were different. I don't mean conceptually they were different. They were sort of the same genre. Just like, say, Chinese and Americans are of the same genre except one speaks Chinese and one speaks English, one lives on one side of the world, one lives on the other side, they go to sleep during your daytime, etc. The details of the two networks were rather different. The ARPANET ran at 50 kilobits per second and the packet radio system ran at 100 or 400 kilobits per second. One had thousand bit uncoded packets; the other had two thousand bit packets which could be coded. The ARPANET assumed that once you sent something it was delivered with a hundred percent reliability. If it didn't get through the system was broken. The other assumed that much of the time you would never get anything through even though the system was working. The protocols that were designed for the ARPANET wouldn't work over the packet radio net because when a packet entered the packet radio net, the only thing the ARPANET would have told it was where it came from but not where it was going. So the packet radio net had no further information to know where to route it. If a packet got lost along the way, the ARPANET hosts would come to a halt. Well, in a radio net you can get interference and so some loss is natural. So we really had to rethink literally the whole issue of host transport protocols. Vint Cerf and I jointly came up with the TCP/IP concept as a new transport mechanism as part of an architecture for internetworking. DARPA then gave a contract to Vint at Stanford to actually implement the TCP/IP concept - along with small efforts at BBN and at University College London. Vint had the lead for developing the specification.

O'NEILL: What were the other two contracts for?

KAHN: That eventually became the IAB. Vint is now the chairman of the IAB. But when we started the internet program in the mid 1970s, originally it was just me in the office running the program. And after Vint was hired, then it was just Vint running the program with me to kibitz. And he was so good at what he did that he basically had everything in his head. What I worried about was what would happen if he got hit by a truck? Number two, what would happen if he would ever have to leave? And number three, how was anybody else in the community ever going to be part of the thinking process.

So he set up, after some discussions, a kind of kitchen cabinet, if you will, of knowledgeable people that he would convene periodically. These were mostly the workers in the field, the key people who were implementing protocols. He would sit down with them and discuss what his plans were and get inputs from them. But it was basically Vint driving that whole process. There were a lot of strong personalities involved, but Vint was so knowledgeable about what was going on, he could have made all the decisions himself, although they gave him a lot of good input. When he left, that group stayed intact. I didn't particularly require it for myself, since Vint and I had created the internet protocols and he and I had it all in our heads before it all started. But in terms of the plan and in the evolution of the research program, it was pretty clear that whether it was Vint or myself or whoever was going to be there in time, you really need inputs from the research community because in the final analysis the problems that needed to be addressed would have to reflect problems that individual groups were having. "Here's a thing they wanted to do; the Internet should help with this. Here's another thing they wanted to do..." So that process became increasingly important as time went on. And, of course, in formulating what the program would be during the year that I had taken it back over again to run it, I relied heavily on the group for inputs.

When I handed it over to Barry Leiner, he had had no direct experience in that area. He had taken over the packet radio program from Vint when Vint left. And when he eventually took over the Internetting program too, he was primed for it. The issues are very unique in that particular discipline. This group was very helpful to him. Barry was the one that reorganized and renamed it the Internet Activities Board, because it had grown

so large by that time. When Vint was involved, the discussions had become so interesting to everybody that the people just wanted to come to hear them. And oftentimes those meetings were held in conjunction with other planned events, like the satellite program that I was running where we would have meetings of people internationally. So he would have one of his ICCB meetings in conjunction with that. Often there would be fifty to a hundred people that would want to sit in as observers just to hear the discussion. Because of that, it almost became unmanageable. When Barry took the program over he said, "I don't want to run a planning group of several hundred people." So he created a smaller body called the IAB thinking that it would remain small. The problem was that pretty soon people who suddenly weren't getting invitations to ICCB meetings because it no longer existed suddenly said "Hey, what happened?" And of course they found out there was an IAB, and then they wanted to come to its meetings. The IAB ended up supporting an Internet Engineering Task Force and an Internet Research Task Force, both reporting to the IAB. The IETF, which Phil Gross chairs, is now a very large open organization. A typical meeting will have hundreds of people. The IRTF is much smaller. So that is sort of how all that evolved. In any event, Barry Leiner took over in 1983, and he ran the program until he left in 1985.

O'NEILL: Okay. That covers the additional questions I had. Do you have any further general statements about networking or ARPA that you would like to add?

KAHN: No.

O'NEILL: Okay. Thank you.

END OF INTERVIEW